

## CLAIMS

What is claimed is:

- 1           1.       A container object stored in platform readable medium executed by a  
2 processor within a platform, the container object comprising:  
3           a hardware identification object to identify to an operating system of the platform that  
4 a type of device represented by the container object is a node; and  
5           a plurality of component objects to identify constituent components of the node.
- 1           2.       The container object of claim 1, wherein the node is a scalability node  
2 controller.
- 1           3.       The container object of claim 2, wherein the container object handles an  
2 ejection notice for a hot-plug removal of the scalability node controller and at least one  
3 component coupled the scalability node controller from the platform.
- 1           4.       The container object of claim 3, wherein the at least one component includes a  
2 processor.
- 1           5.       The container object of claim 4, wherein the at least one component includes a  
2 memory.
- 1           6.       The container object of claim 3, wherein the ejection notice is provided by an  
2 operating system of the platform.
- 1           7.       The container object of claim 1, wherein one of the plurality of component  
2 objects includes a processor object to identify one or more processors coupled to the node.
- 1           8.       The container object of claim 7, wherein one of the plurality of component  
2 objects includes a device object to identify one or more memory devices coupled to the node.

1           9.       The container object of claim 8, wherein any hotplug operation is applied as a  
2 single operation to the container object and its constituent objects.

1           10.      The container object of claim 8, wherein hotplug addition of the one or more  
2 processors and the one or more memory devices belonging to the container object is notified  
3 to the operating system simultaneously in a single operation.

1           11.      The container object of claim 8, wherein the hotplug removal of the processors  
2 and the memory devices belonging to the container object is notified to the operating system  
3 simultaneously in a single operation.

1           12.      The container object of claim 8, wherein the container object handles a hot-  
2 plug addition by activating the memory devices associated with the device object online prior  
3 to activating the processors associated with the processor object.

1           13.      The container object of claim 1 further comprising a proximity object to  
2 describe a proximity domain within the platform that the node belongs to.

1           14.      The container object of claim 10, wherein the proximity object is an assigned  
2 integer value to represent the proximity domain that the node belongs to.

1           15.      A platform comprising:  
2           a memory device to store a system Basic Input/Output System (BIOS); and  
3           a plurality of processor substrates in communication with the memory device, the  
4 plurality of processor substrates including  
5           a first processor substrate including a first plurality of components and a first  
6 storage device to contain a first initialization BIOS to initialize the first plurality of  
7 components in response to hot-plug addition of the first processor substrate to the  
8 platform, and  
9           a second processor substrate including a second plurality of components and a  
10 second storage device to contain a second initialization BIOS to initialize the second  
11 plurality of components in response to hot-plug addition of the second processor  
12 substrate to the platform.

1 16. The platform of claim 12, wherein the first plurality of components of the first  
2 processor substrate include a processor cluster and a memory cluster.

1 17. The platform of claim 13, wherein the first plurality of components of the first  
2 processor substrate further include a scalability node controller coupled to the processor  
3 cluster and the memory cluster.

1 18. The platform of claim 14, wherein the scalability node controller of the first  
2 plurality of components is coupled to local memory of the memory cluster through a plurality  
3 of communication sub-links supporting a total data throughput of at least one Gigabyte per  
4 second.

1 19. The platform of claim 14, wherein the scalability node controller of the first  
2 plurality of components includes a plurality of scalability port interfaces coupled to a first  
3 connector of the first processor substrate.

1 20. The platform of claim 16, wherein the second plurality of components of the  
2 second processor substrate further include a secondary scalability node controller including a  
3 plurality of scalability port interfaces coupled to a second connector of the second processor  
4 substrate.

1 21. The platform of claim 17 further comprising:

2 an interconnect substrate including a third connector, a fourth connector and a fifth  
3 connector, the third connector adapted to mate with the first connector of the first processor  
4 substrate and coupled to the fifth connector via a first link and a second link, the fourth  
5 connector adapted to mate with the second connector of the second processor substrate and  
6 coupled to the fifth connector via a third link and a fourth link.

1 22. The platform of claim 18 further comprising:

2 an input/output (I/O) substrate including a sixth connector coupled to the fifth  
3 connector, the I/O substrate further includes (1) a first scalability port switch coupled to the  
4 first link and the third link, (2) a second scalability port switch coupled to the second link and  
5 the fourth link, (3) a first Server Input/Output Hub coupled to the first and second scalability

6 port switches, and (4) a second Server Input/Output Hub coupled to the first and second  
7 scalability port switches.

23. The platform of claim 19, wherein the first initialization BIOS contained in the first storage device, when executed, also establishes a communication path between the scalability node controller and at least one of the first scalability port switch and the second scalability port switch.

24. The platform of claim 19, wherein the second initialization BIOS contained in the second storage device, when executed, also establishes a communication path between the secondary scalability node controller and at least one of the first scalability port switch and the second scalability port switch.

0901 25. The platform of claim 13, wherein the first initialization BIOS contained in the  
0902 first storage device elects a processor from a plurality of processors associated with the  
0903 processor cluster to act as a node boot strap processor for the first processor substrate.

26. The platform of claim 12, wherein the second initialization BIOS contained in the second storage device elects a processor from a plurality of processors implemented on the second processor substrate to act as a node boot strap processor for the second processor substrate.

1            27.     The platform of claim 12, wherein the first initialization BIOS contained in the  
2     first storage device initializing the first plurality of components prior to notification of an  
3     operating system running on the platform of the initialized hot-plugged first processor  
4     substrate.

1           28.     A platform comprising:  
2           a first processor substrate including a first plurality of components and a first storage  
3     device to contain a first code segment of Basic Input/Output System (BIOS) that, when  
4     executed, initializing the first plurality of components in response to hot-plug addition of the  
5     first processor substrate to the platform; and  
6           a second processor substrate including a second plurality of components and a second  
7     storage device to contain a second code segment of BIOS that, when executed, initializing the

8 second plurality of components in response to hot-plug addition of the second processor  
9 substrate to the platform.

1 29. The platform of claim 25 further comprising:

2 an interconnect substrate including a first connector, a second connector and a third  
3 connector, the first connector being adapted to mate with a connector of the first processor  
4 substrate that is coupled to a plurality of scalability port interfaces of a first scalability node  
5 controller of the first plurality of components, the second connector being adapted to mate  
6 with a connector of the second processor substrate that is coupled to a plurality of scalability  
7 port interfaces of a second scalability node controller of the second plurality of components,  
8 and a third connector coupled to both the first connector via a first link and a second link and  
9 the second connector via a third link and a fourth link.

1 30. The platform of claim 26 further comprising:

2 an input/output (I/O) substrate including a fourth connector coupled to the third  
3 connector, the I/O substrate further includes (1) a first scalability port switch coupled to the  
4 first link and the third link, and (2) a second scalability port switch coupled to the second link  
5 and the fourth link.

1 31. The platform of claim 27, wherein the first code segment of BIOS, when  
2 executed, also establishes a communication path between the first scalability node controller  
3 and at least one of the first scalability port switch and the second scalability port switch.

1 32. The platform of claim 28, wherein the second code segment of BIOS, when  
2 executed, also establishes a communication path between the second scalability node  
3 controller and at least one of the first scalability port switch and the second scalability port  
4 switch.

1 33. The platform of claim 25, wherein the first code segment of BIOS initializing  
2 the first plurality of components prior to notification of an operating system running on the  
3 platform of the first plurality of components upon hot-plug addition of the first processor  
4 substrate.

1 34. A platform comprising:

an operating system;  
an input/output (I/O) substrate including (i) a first scalability port switch, (ii) a second scalability port switch, (iii) a first server I/O hub and (iv) a second server I/O hub;  
a first processor substrate including a first plurality of components in communication with the first and second scalability port switches; and  
a second processor substrate including a second plurality of components and circuitry to support dynamic partitioning of the platform by signaling the operating system of a hot-plug removal of the second server I/O hub and the second plurality of components to cause the operating system to configure the first scalability port switch and the second scalability port switch so as to partition the platform into a first platform including the first plurality of components and the first server I/O hub and a second platform including the second plurality of components and the second server I/O hub.

35. The platform of claim 34, wherein the second platform is able to run an updated software application independent of and without interrupting operations of the first platform.

36. The platform of claim 35, wherein the updated software application is an updated operating system.

37. A method comprising:  
providing a multi-node platform under control of an operating system, the multi-node platform including a first processor substrate and a second processor substrate in communication with an input/output (I/O) substrate; and  
implementing a portion of a Basic Input/Output Subsystem (BIOS) on the first processor substrate to initialize components on the first processor substrate in response to hot-plug addition of the first processor substrate before joining the running operating system.

38. The method of claim 37 further comprising:  
implementing a portion of the BIOS on the second processor substrate to initialize components on the second processor substrate in response to hot-plug addition of the second processor substrate before joining the running operating system.

39. The method of claim 37, wherein the components on the first processor substrate include at least two processors.